



(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number:

0 373 361
A2

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 89120667.4

(51) Int. Cl.⁵: G06F 9/45

(22) Date of filing: 08.11.89

(33) Priority: 16.12.88 US 285459

(71) Applicant: International Business Machines Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)

(43) Date of publication of application:
20.06.90 Bulletin 90/25

(72) Inventor: Hopkins, Martin Edward
300 Douglas Road
Chappaqua New York 10514(US)
Inventor: Warren, Henry Stanley, Jr.
Kitchawan Road
Ossining New York 10562(US)

(84) Designated Contracting States:
DE FR GB

(74) Representative: Jost, Ottokar, Dipl.-Ing.
IBM Deutschland GmbH. Patentwesen und
Urheberrecht Schönaicher Strasse 220
D-7030 Böblingen(DE)

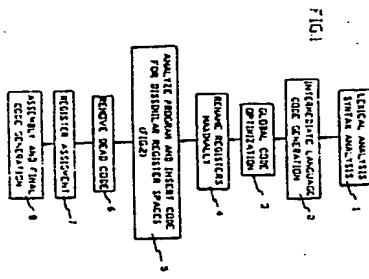
(54) Generating efficient code for a computer with dissimilar register spaces.

(57) A compiler generates compiled object code from source code of a computer program in a manner that produces efficient object code for a computer with dissimilar register spaces. The technique comprising the steps of

generating code that references symbolic registers in which the register class is not distinguished, making entries in a table denoting the context in which each symbolic register occurs and constructing an equivalence tree of symbolic registers for move instructions assigned to a same equivalence class, for each equivalence class, forming the logical OR function of register usage information for all symbolic registers in the class, and for each symbolic register that appears in more than one register space context, generating new symbolic register numbers so that there is one number for each register space, and storing the numbers in said table, and

if a definition point of a symbolic register is encountered and that symbolic register is used in more than one register space context, inserting code in said program to either do the same operation as is done at the definition point in each register space or move a value in the symbolic register from one space to another. The improvement achieved is in object code space and time of execution.

A2
EP 0 373 361



Xerox Copy Centre

GENERATING EFFICIENT CODE FOR A COMPUTER WITH DISSIMILAR REGISTER SPACES

This invention generally relates to computer compilers and, more particularly, to a method that permits an optimizing compiler to generate more efficient code for computers that have dissimilar register spaces. The invention has specific application to optimizing compilers for reduced instruction set computers (RISC) but could be used with advantage in any compiler.

5 There currently exists a need for a technique that may be used to improve the quality of the object code produced by a compiler. The technique should apply to a machine that has multiple register spaces, with the properties that some operations can only be done in one of the spaces, some can be done in either, and it is time-consuming to move data from one space to another. Examples of register spaces are general purpose registers (GPRs), floating point registers (FPRs) and condition registers (CRs). As an 10 example of a time-consuming move, on the IBM Systems 370 (S/370) series of computers, and in many other computers, the only way to move data between a GPR and an FPR is to go through storage; i.e., to execute a Store Float followed by a Load or a Store followed by a Load Float.

It is therefore an object of this invention to provide a technique for improving the quality of the object code produced by a compiler.

15 It is another object of the invention to provide an improvement in both the object code space and time to execute the code.

It is a further object of the invention to provide a technique for generating efficient code for a computer with dissimilar register spaces.

According to the invention, a compiler generates compiled object code from source code of a computer 20 program in a manner that produces efficient object code for a computer with dissimilar register spaces. The technique comprising the steps of (1) generating code that references symbolic registers in which the register class is not distinguished, (2) making entries in a table denoting the context in which each symbolic register occurs and constructing an equivalence tree of symbolic registers for "move" instructions assigned to a same equivalence class, (3) for each equivalence class, forming the logical OR function of register 25 usage information for all symbolic registers in the class, and for each symbolic register that appears in more than one register space context, generating new symbolic register numbers so that there is one number for each register space, and storing the numbers in said table, and (4) if a definition point of a symbolic register is encountered and that symbolic register is used in more than one register space context, inserting code in said program to either do the same operation as is done at the definition point in 30 each register space or move a value in the symbolic register from one space to another. The improvement achieved is in object code space and time of execution.

The foregoing and other objects, aspects and advantages of the invention will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

35 Fig. 1 is a flow diagram showing the overall structure of a typical optimizing compiler illustrating where the present invention fits into this structure;

Fig. 2 is a flow diagram illustrating the logic of the present invention;

Fig. 3 is a detailed flow diagram of the pass 1 logic shown in Fig. 2;

Fig. 4 is a detailed flow diagram of the interlude logic shown in Fig. 2; and

40 Fig. 5 is a detailed flow diagram of the pass 2 logic shown in Fig. 2.

Referring now to the drawings, and more particularly to Fig. 1, there is shown a high level flow diagram of an optimizing compiler. As is typical of such compilers, the process begins at block 1 with a lexical analysis and syntax analysis of the source code being compiled. Next, in block 2, an intermediate language code is generated, and this is followed in block 3 by a global code optimization. The process described thus far is conventional and well understood by those skilled in the design of compilers.

45 The improvement according to the present invention begins with block 4 where the registers are renamed maximally. Next, the program is analyzed in block 5 and code is inserted for dissimilar register spaces. This process is illustrated in more detail in Fig. 2 to which reference is now made. In block 11, a pass is made over the program to collect register usage and equivalence class information. This information is then processed in block 12 by forming the logical OR of the register usage information for all the registers in each equivalence class. Once this is done, a second pass over the program is made in block 13. In this pass, code is inserted in the program for the multiple register spaces. Thus, from Fig. 2, it will be appreciated that two passes are made over the program, but that the second pass is separated from the first by an interlude for processing the register usage and equivalence class information gathered in the first pass.

Returning now to Fig. 1 of the drawings, when the processing of Fig. 2 is completed, dead code in the program is removed in block 6. This is desirable because the steps of Fig. 2, block 13, may have inserted some unnecessary (i.e., "dead") code, and a simple way to detect and remove it is to perform the well-known optimization of dead code removal at this point. Register assignments are made in block 7 on the basis of the code as it exists to that point. Finally, the assembly and final code generation is performed in block 8 to complete the process.

Further elaborating, the technique according to the invention is embodied in the following steps:

1. First, code is generated that references symbolic registers, in which the register class is not distinguished. That is, a symbolic register such as R100 might later become assigned to a GPR, to an FPR or to a CR, depending on how it is used. This is generally illustrated at block 2 in Fig. 1. The usual optimizations such as code motion and commoning are done on this code, in block 3. The code at this point should have distinct symbolic register numbers whenever possible (a concept called the "maximal number of names"), so that step (4) below will only insert extra instructions when there is a high probability that they are necessary. This is an observation affecting only compile time, not object code quality.

2. Next, as illustrated in block 11 in Fig. 2, a pass is made over the program, noting (in a table) the context in which each symbolic register (SR) appears. This is a simple linear pass, ignoring the branching structure and the definition-use relation. During this pass, move register (MR) instructions are treated specially. The instruction "MR rx,ry" does not denote any particular register space for rx or ry, but it does mean that they are in the same spaces. To handle this, when "MR rx,ry" is encountered, we put rx and ry in the same equivalence class. Any standard algorithm for defining equivalence classes of the entries in a table may be used, for example Algorithm E of Knuth, The Art of Computer Programming, Fundamental Algorithms, Vol. 1, section 2.3.3 on page 354. This simple algorithm constructs a tree for each equivalence class, by putting a pointer from rx to ry in the table.

3. Next, the processing of block 12 in Fig. 2 is made by a pass over the table. For each equivalence class, the logical OR of the register usage information for all SRs in the class is formed. Then, for each SR that appears in more than one register space context, we generate new SR numbers, so that there is one for each register space. These new SR numbers are saved in the table.

4. Next, the second linear pass shown in block 13 in Fig. 2 is made over the program in which "fix up" code is inserted. If a definition point of an SR is encountered, and that SR is used (anywhere in the programs) in more than one register space context, then code is inserted to either (a) do the same operation as is done at the definition point, in each register space, or (b) move the value from one space to another. During this second pass, if a "use" point of an SR is encountered, and that SR appears in more than one context, then the SR of the use is changed to be the appropriate SR number for the context of the use.

5. Lastly, the standard optimization of "dead code elimination" is done as shown in block 6 of Fig. 1, and then register allocation, as shown in block 7 in Fig. 1. Register allocation is run separately on each register space.

Steps 2 to 4 above form this invention, whereas steps 1 and 5 are a common part of an optimizing compiler. The "fix up" code could be inserted either at definition points or at use points. Inserting it at definition points (as described above) is usually best, because often one definition point feeds several uses, and because sometimes the definition is a simple one, such as setting a variable equal to a constant, that can be done in any of the register spaces (avoiding a "move" from one space to another).

The logic of pass 1, represented by block 11 in Fig. 2, is shown in more detail in Fig. 3, to which reference is now made. The process begins by retrieving the first or next instruction in function block 15. A test is made in decision block 16 to determine if the instruction is fake. If it is, the process jumps to decision block 25; otherwise, a test is made in decision block 17 to determine if the instruction is a move register (MR) instruction. If it is, the source and target registers are put in the same equivalence class by making the higher numbered one point to the lower numbered one in function block 18. The process then jumps to decision block 25. If, on the other hand, the test in decision block 17 is negative, a test is made in decision block 19 to determine if the instruction is a conditional branch instruction. If it is, the "CR context" for the tested register is noted in function block 20, and then the process jumps to decision block 25. If the instruction is not a conditional branch instruction, a test is made in decision block 21 to determine if the instruction is a compare instruction. If so, the "CR context" for the target register is noted in function block 22. In either case, the "FPR context" is noted for all FPRs referenced in function block 23. Next, the "GPR context" is noted for all GPRs referenced in function block 24. Finally, in decision block 25 a test is made to determine if there are any more operations to be processed. If there are, the process loops back to function block 15; otherwise, control passes to the interlude logic, shown in Fig. 4.

In Fig. 4, the interlude process begins by initializing in function block 27; that is, i is set equal to the

first register. Then a test is made in decision block 28 to determine if i is less than or equal to the last register. If it is, a second test is made in decision block 29 to determine if i is in any equivalence class. If not, i is put in a class by itself in function block 30. In either case, j is set equal to the next in class after i, and k is set equal to next in class after j in function block 31. In function block 32, the usage of k is set equal to the logical OR of the usage of k and the usage of i, and in function block 33, i is made to point to k. Then, in function block 34, i is indexed by adding one to i before looping back to decision block 28.

When the result of the test in decision block 28 is negative, the process goes to function block 36 where, again, i is set equal to the first register. A test is made in decision block 37 to determine if i is less than or equal to the last register. If so, a further test is made in decision block 38 to determine if i appears in more than one context. If so, new names are made up in function block 39 so that there is a different name for each context. In either case, i is then indexed in function block 40 by adding one to i. The process ends when the result of the test in decision block 37 is negative, at which point control passes to the logic of pass 2 shown in Fig. 5.

Processing for pass 2 begins by retrieving the first or next instruction in function block 42. A test is made in decision block 43 to determine if the instruction is fake. If so, the process jumps to decision block 52; otherwise, a test is made in decision block 44 to determine if the instruction is a move register (MR) instruction. If it is, in function block 45 code is inserted for each register space in which the target appears to do the "move" instruction, and then the process jumps to decision block 52. If not a move register instruction, a test is made in decision block 46 to determine if the target register appears in more than one context. If so, a further test is made in decision block 47 to determine if the operation can be done efficiently in the alternate register spaces. If so, code is inserted to do the operation in the alternate register spaces in function block 48. If not, code is inserted to move the result to the alternate register spaces in function block 49. In either case or if the result of the test in decision block 46 is negative, control goes to decision block 50 where a test is made to determine if any source register appears in more than one context. If so, the register name is changed to the new name that is in the appropriate register space for the context. If so, the register name is changed to the new name that is in the appropriate register space for the context. In either case, control next goes to decision block 52 where a test is made to determine if there are any more operations to be processed. If so, the process loops back to function block 42; otherwise, the process ends.

As an example of these transformations, assume that the computer has multiple general registers and multiple condition registers. Assume for simplicity that the condition registers are the same length (e.g., 32 bits) as the general registers. This invention is applicable to the case in which the registers are of different lengths; the details are more complicated, but the principles are the same. Then some typical transformations are shown below. Here, "rx" in the left column denotes the original symbolic register. We assume that "rx" appears in both contexts (general register and condition register) somewhere in the program. In the right column, "crx" denotes the condition register counterpart of "rx". That is, when registers are assigned to the code in the right column, "rx" will be assigned to a GPR, and CRX will be assigned to a condition register.

40	L	rx,...	L	rx,...	(1)
			MTCR	crx,rx	
	LI	rx,imm	LI	rx,imm	(2)
			LICR	crx,imm	
45	MR	rx,ry	MR	rx,ry	(3)
			MCR	crx,cry	
	CMP	rx,...	CMP	crx,...	(4)
50	BGT	rx,label	BGT	crx,label	(5)

(1) shows a fixed point load instruction. Since rx appears somewhere in both contexts (we are assuming), we generate a new symbolic register crx and insert a "move to condition register" instruction to move rx to crx. MTCR is a slow, undesirable instruction, but we are assuming that there is no way to load directly from memory to a CR, so we had to use it.

(2) shows a Load Immediate instruction. We assume the computer has a fast Load Immediate into a CR (LICR), so we use that to set crx, rather than an MTCR.

(3) shows a Move Register instruction. If either rx or ry appears in both contexts, then they both will appear in both contexts, by the equivalence class algorithm described above. Hence, both rx and ry have CR counterparts crx and cry. We insert an instruction to move cry to crx.

- 5 (4) shows a Compare instruction. We assume the target can only be a CR. Hence, we change rx to the CR counterpart crx.

- 10 (5) shows a Branch on Condition. We assume the tested register rx can only be a CR, so we change it to the CR counterpart crx.

15 Assume a program contains the source code shown in column (1) below, written in a high level language such as PL/I. Variable X is a logical variable (a single bit). We do not show any flow of control; the lines can be anywhere in the program. Some may be in loops and they are not necessarily in the order shown. A compiler that does not incorporate the present invention would generate code similar to that shown in column (2). It contains an MTCR (move to condition register) instruction which we assume is costly to execute. A compiler that does incorporate the present invention would generate the code similar to that shown in column (3). This code is faster because it employs LICR (load immediate CR) instructions, rather than an MTCR instruction.

	(1)	(2)		(3)	
20	X=0;	LI	R1,0	LI LICR	R1,0 CR1,0
	X=1;	LI	R1,1	LI LICR	R1,1 CR1,1
25	IF X THEN GO TO LAB;	MTCR BC	CR1,R1 CR1,LAB	BC	CR1,LAB
	Y=X;	ST	R1,Y	ST	R1,Y

- 30 If the Store into Y were not present, then the code in column (3) would simplify; that is, GPR R1 would become dead (have no uses), so the conventional optimization of dead code elimination would eliminate the two LI instructions. But the code in column (2) does not simplify in this case.

- 35 While the invention has been described in terms of a single preferred embodiment, those skilled in the art will recognize that the invention may be practiced with modification within the spirit and scope of the appended claims.

Claims

- 40 1. In a compiler which generates compiled object code from source code of a computer program, a computer performed method of generating efficient object code for a computer with dissimilar register spaces comprising the steps of:
generating by a computer code that references symbolic registers in which the register class is not distinguished (2);
- 45 making entries in a table by the computer denoting the context in which each symbolic register occurs and constructing an equivalence tree of symbolic registers for move instructions assigned to a same equivalence class (11);
for each equivalence class, forming by the computer the logical OR function of register usage information for all symbolic registers in the class, and for each symbolic register that appears in more than one register
- 50 space context, generating by the computer new symbolic register numbers so that there is one number for each register space, and storing the numbers in said table (12); and if a definition point of a symbolic register is encountered and that symbolic register is used in more than one register space context, inserting by the computer code in said program to either do the same operation as is done at the definition point in each register space or move a value in the symbolic register from one space to another (13).

55 2. The method of generating efficient object code for a computer as recited in claim 1 wherein said step of making entries in a table (11) further comprises the step of checking by the computer each instruction to determine if it is a move register instruction (17), and if so, putting source and target registers in a same

equivalence class (18).

3. The method of generating efficient object code for a computer as recited in claim 1 or 2 wherein said step of making entries in a table (11) further comprises the step of checking by the computer each instruction to determine if it is a conditional branch instruction (19), and if so, noting a conditional register context for that register (20).

5 4. The method of generating efficient object code for a computer as recited in one of claims 1 to 3 wherein said step of making entries in a table (11) further comprises the step of checking by the computer each instruction to determine if it is a compare instruction (21), and if so, noting a conditional register context for that register (22).

10 5. The method of generating efficient object code for a computer as recited in one of claims 1 to 4 wherein the step of making entries in a table (11) further comprises the steps of:
noting by the computer a floating point register context for all referenced floating point registers (23); and
noting by the computer a general purpose register context for all referenced general purpose registers (24).

15 6. The method of generating efficient object code for a computer as recited in one of claims 1 to 5 wherein the steps of forming the logical OR function and generating new symbolic register numbers (12) further comprise the steps of:
testing by the computer each register to determine if it is assigned to an equivalence class (unreferenced block left of 30), and if not, putting that register in a: class by itself (30) and
testing by the computer each register to determine if it appears in more than one context (38), and if so,
20 making up new names so that there is a different one for each context (39).

25 7. The method of generating efficient object code for a computer as recited in one of claims 1 to 6 wherein said step of inserting code (13) further comprises the step of checking by the computer each instruction to determine if it is a move register instruction (44), and if it is, inserting by the computer code to do the move instruction in that register space (45).

30 8. The method of generating efficient object code for a computer as recited in one of claims 1 to 7 wherein said step of inserting code (13) further comprises the step of checking by the computer each target register to determine if that register appears in more than one context (46), and if so, determining by the computer if the operation can be done efficiently in alternate register spaces (47), and if so, inserting by the computer code to do the operation in the alternate register spaces (48), but otherwise, inserting by the
computer code to move the result to the alternate register spaces (49).

35 9. The method of generating efficient object code for a computer as recited in one of claims 1 to 8 wherein said step of inserting code (13) further comprises the steps of testing by the computer each source register to determine if it appears in more than one context (50), and if so, changing by the computer the register to a new name that is in the appropriate space for a corresponding operation (51).

35

40

45

50

55

FIG.1

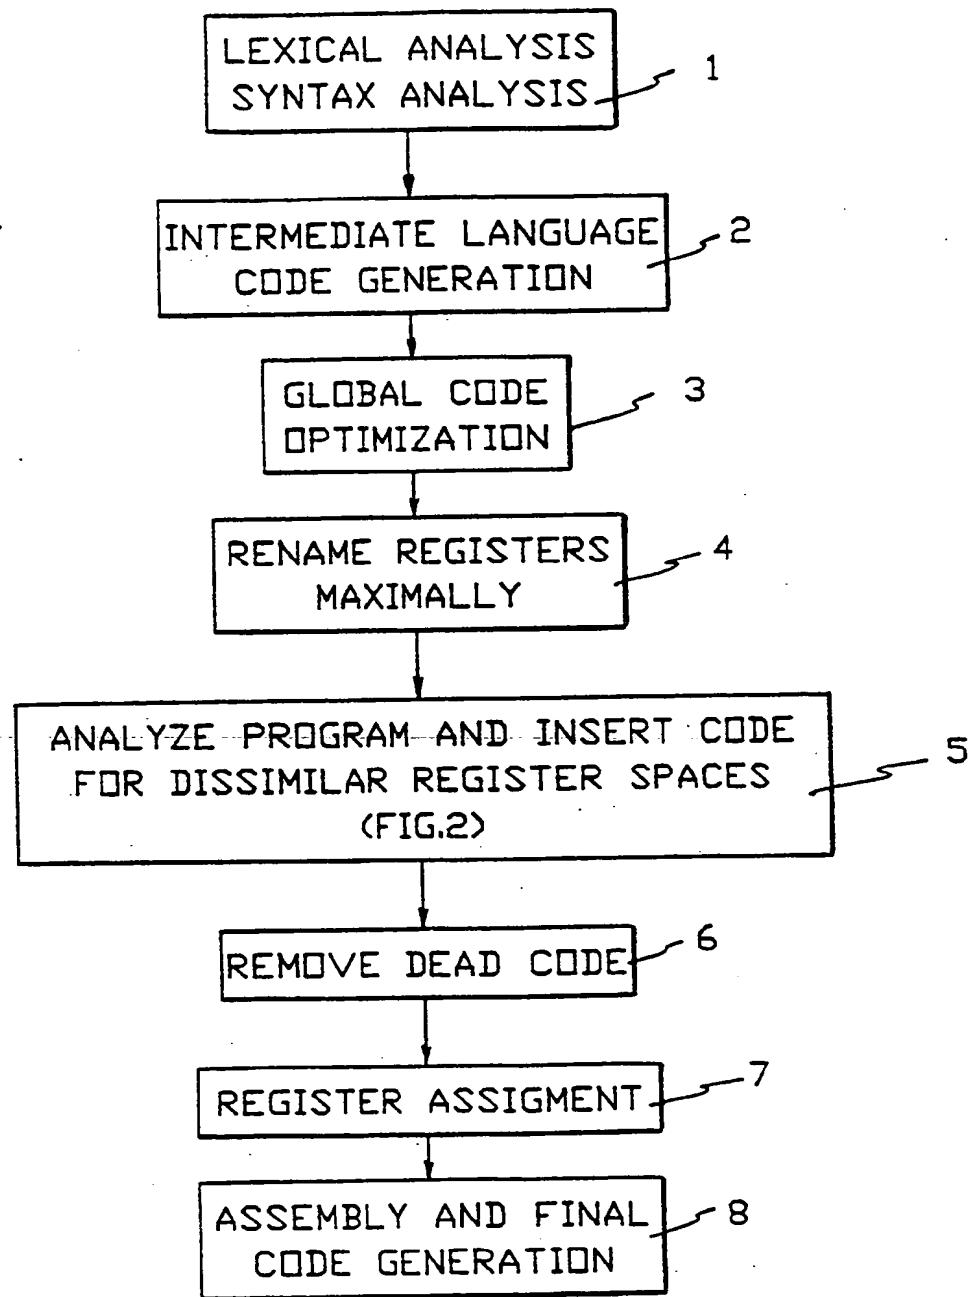


FIG.2

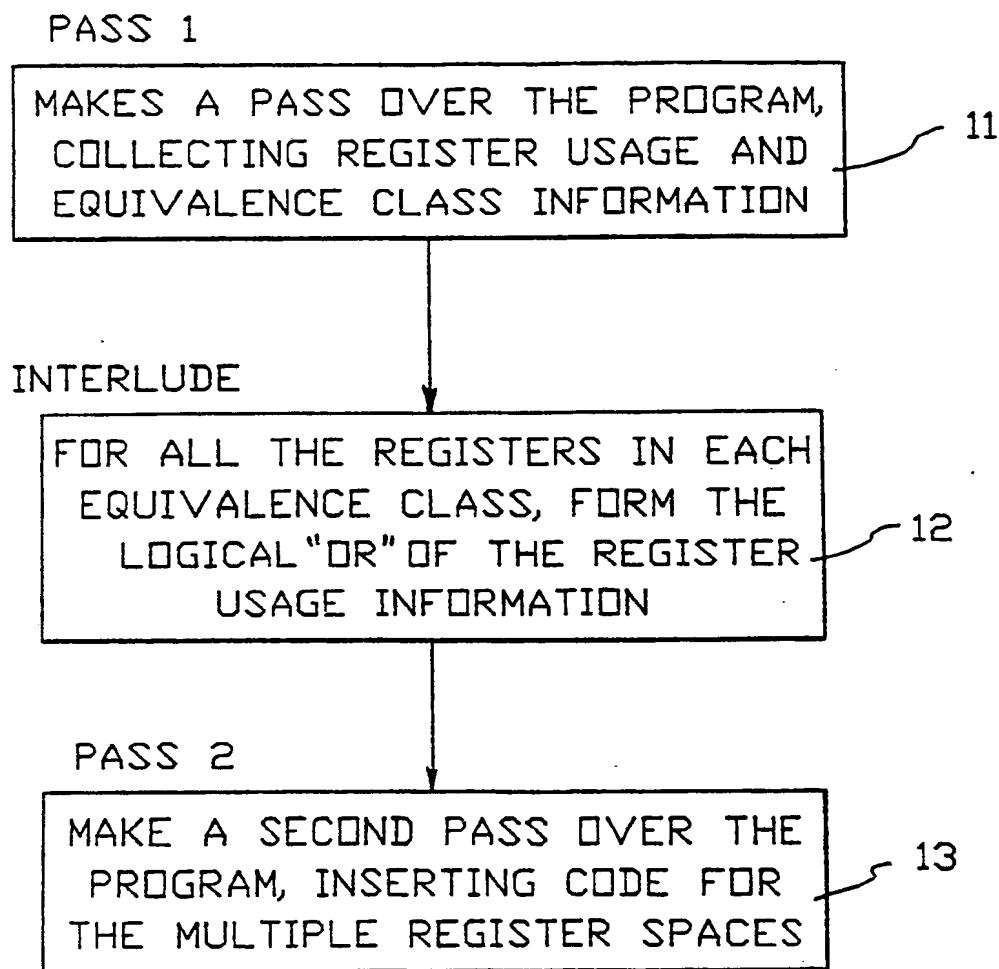
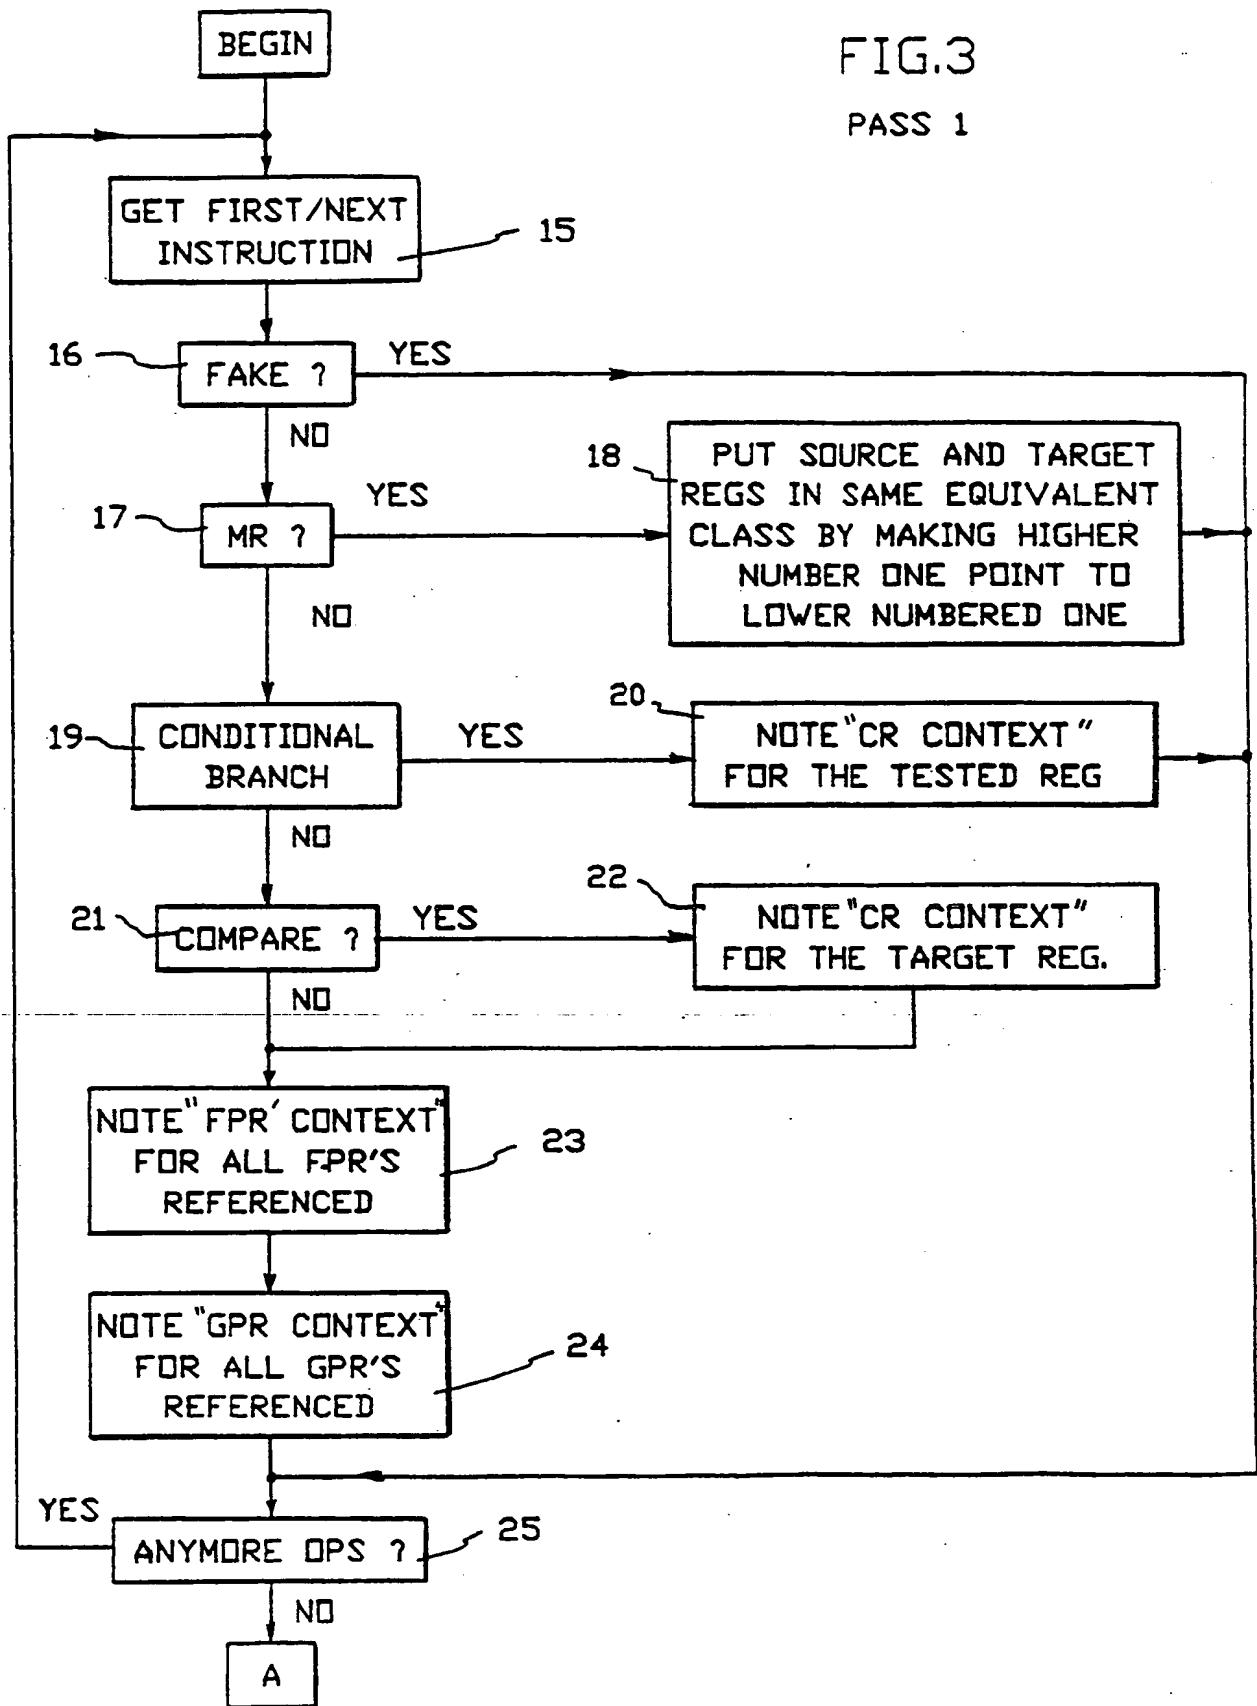


FIG.3

PASS 1



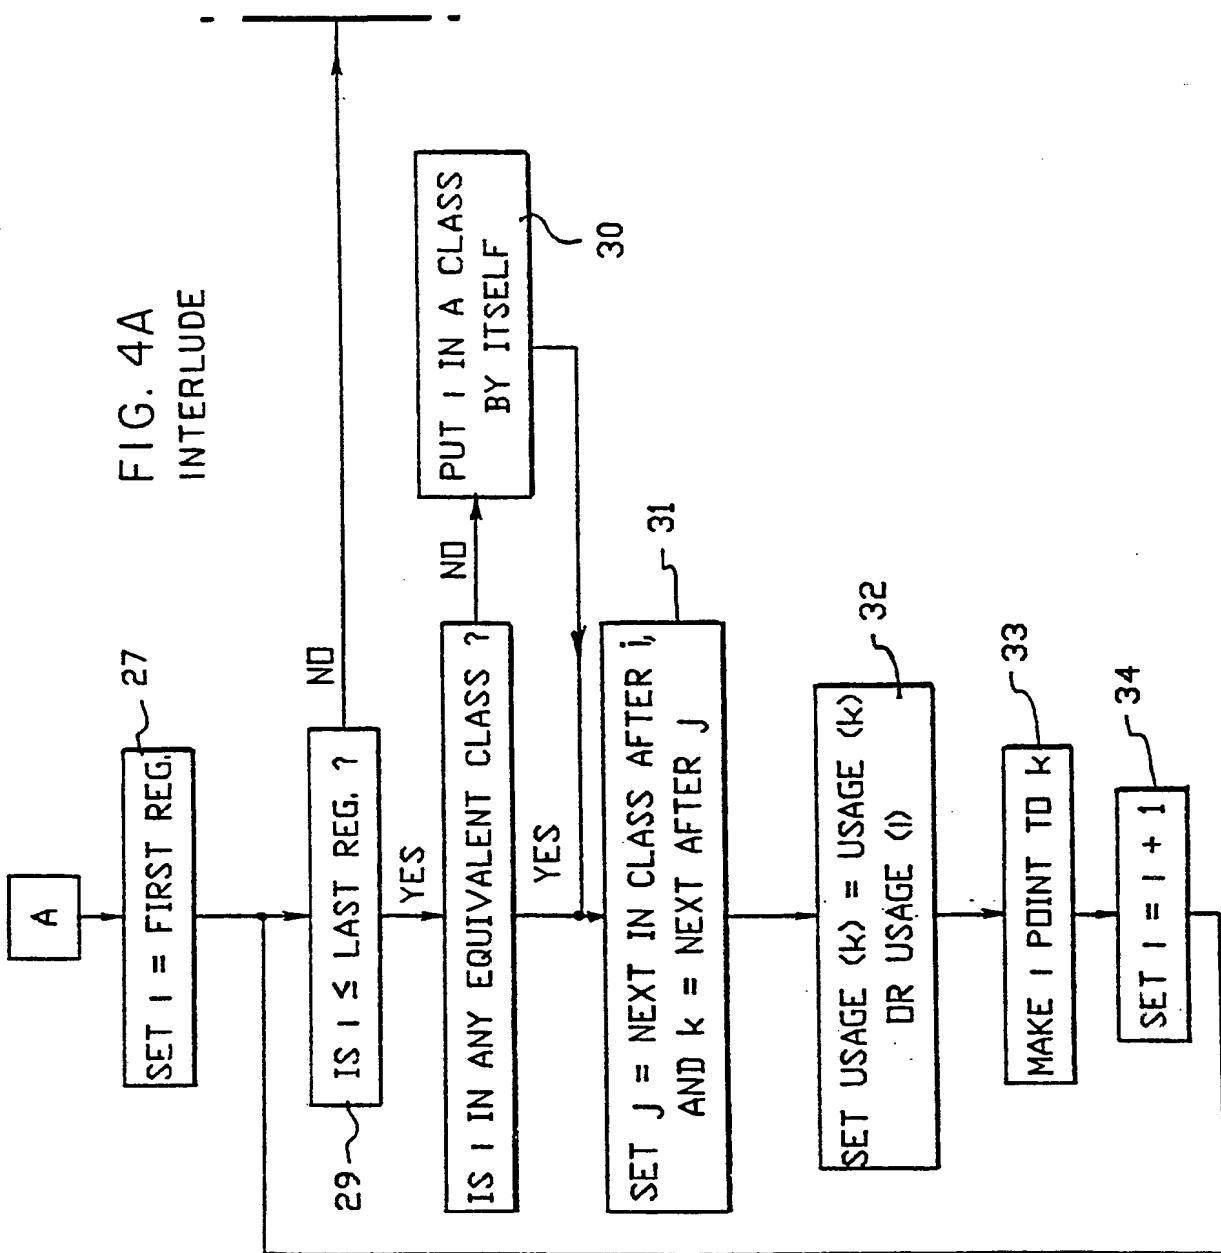


FIG. 4B
INTERLUDE

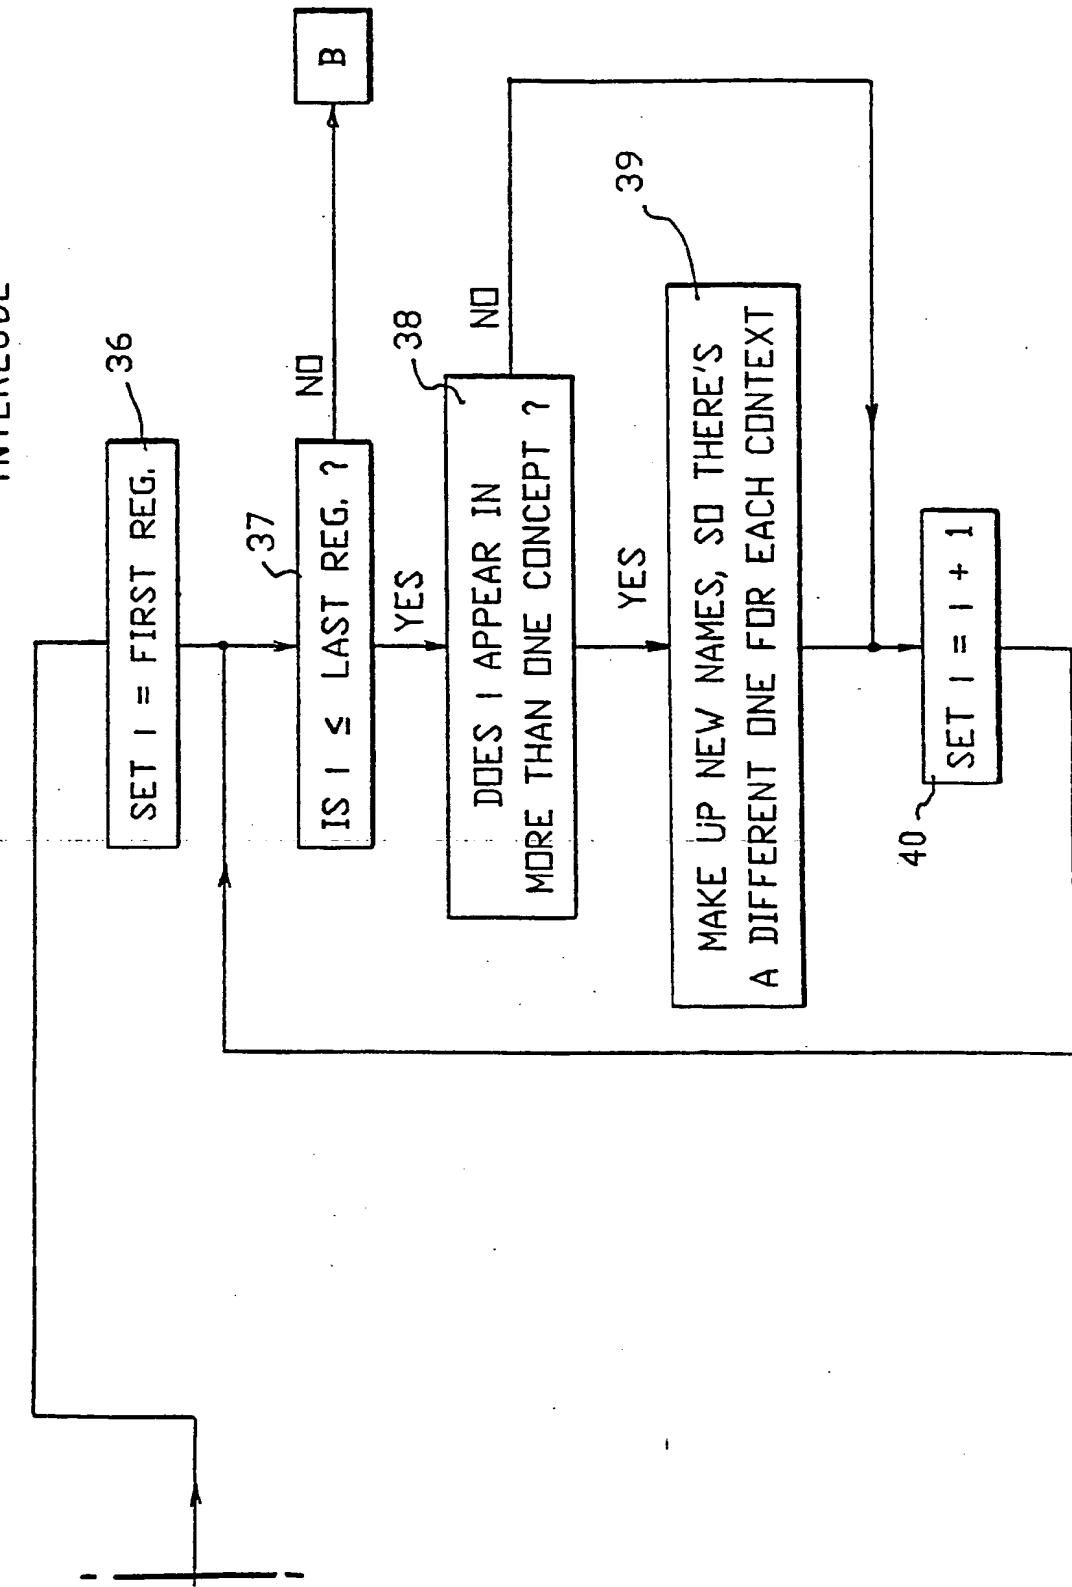
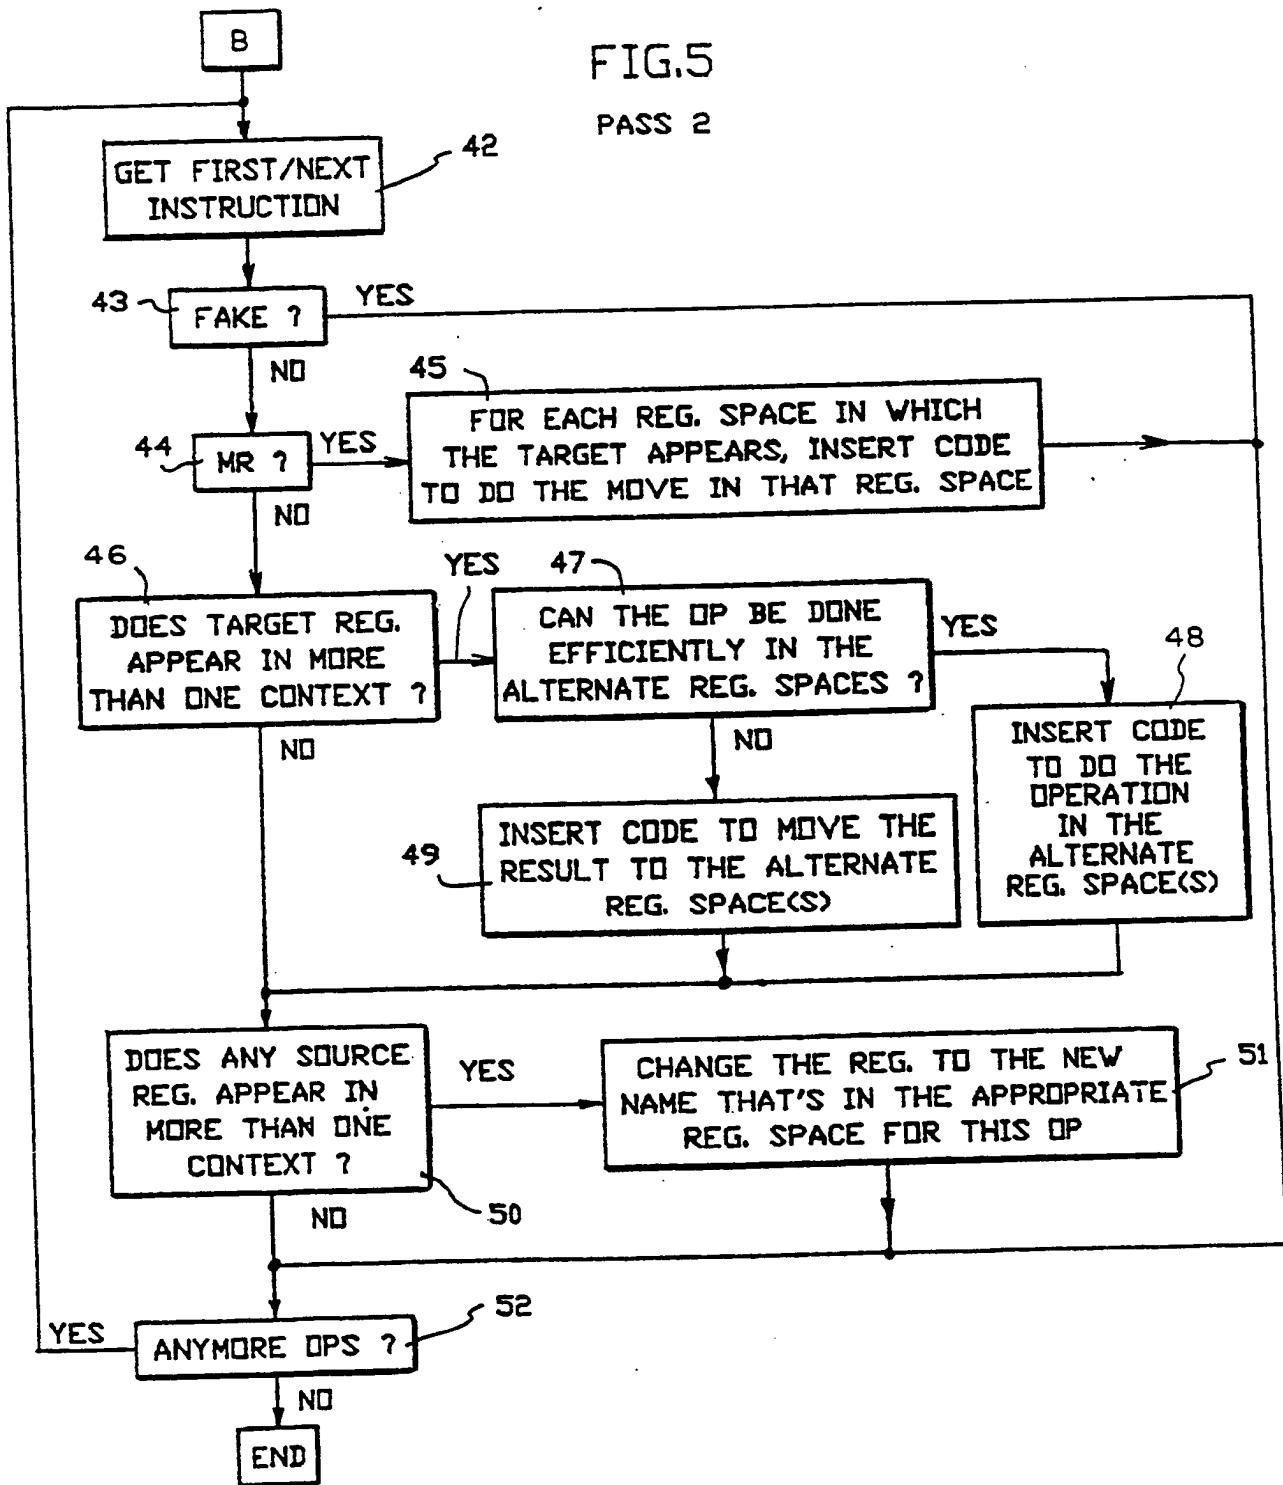


FIG.5

PASS 2





(19) Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 373 361 A3

(12) **EUROPEAN PATENT APPLICATION**

(21) Application number: 89120667.4

(51) Int. Cl.5: G06F 9/45

(22) Date of filing: 08.11.89

(30) Priority: 16.12.88 US 285459

Armonk, N.Y. 10504(US)

(43) Date of publication of application:
20.06.90 Bulletin 90/25

(72) Inventor: Hopkins, Martin Edward
300 Douglas Road
Chappaqua New York 10514(US)
Inventor: Warren, Henry Stanley, Jr.
Kitchawan Road
Ossining New York 10562(US)

(60) Date of deferred publication of the search report:
29.01.92 Bulletin 92/05

(74) Representative: Jost, Ottokar, Dipl.-Ing.
IBM Deutschland GmbH Patentwesen und
Urheberrecht Schönaicher Strasse 220
W-7030 Böblingen(DE)

(71) Applicant: International Business Machines
Corporation
Old Orchard Road

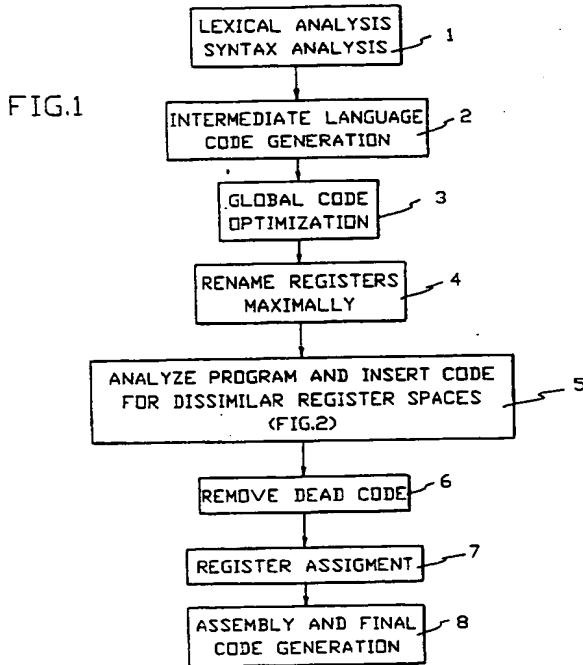
(54) Generating efficient code for a computer with dissimilar register spaces.

(57) A compiler generates compiled object code from source code of a computer program in a manner that produces efficient object code for a computer with dissimilar register spaces. The technique comprising the steps of

generating code that references symbolic registers in which the register class is not distinguished, making entries in a table denoting the context in which each symbolic register occurs and constructing an equivalence tree of symbolic registers for move instructions assigned to a same equivalence class,

for each equivalence class, forming the logical OR function of register usage information for all symbolic registers in the class, and for each symbolic register that appears in more than one register space context, generating new symbolic register numbers so that there is one number for each register space, and storing the numbers in said table, and if a definition point of a symbolic register is encountered and that symbolic register is used in more than one register space context, inserting code in said program to either do the same operation as is done at the definition point in each register space or move a value in the symbolic register from one space to another. The improvement achieved is in object

code space and time of execution.



EP 0 373 361 A3



European
Patent Office

EUROPEAN SEARCH
REPORT

Application Number

EP 89 12 0667

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
A	A.V. AHO ET AL. 'COMPILERS PRINCIPLES, TECHNIQUES, AND TOOLS' 1987 , ADDISON-WESLEY PUBLISHING CO. , US * page 541, last paragraph - page 546, line 31 ** - - -	1	G 06 F 9/45
A	PROCEEDINGS OF THE SIGPLAN'82 SYMPOSIUM ON COMPILER CONSTRUCTION June 1982, BOSTON, US pages 22 - 31; M. AUSLANDER ET AL.: 'An overview of the PL8 compiler' - - -	1	
A	IBM TECHNICAL DISCLOSURE BULLETIN. vol. 28, no. 10, March 1986, NEW YORK US pages 4506 - 4513; 'Register allocation' * page 4506, paragraph 2 *** page 4509, last paragraph ** - - -	1	
A	PROCEEDINGS OF THE SIGPLAN '82 SYMPOSIUM ON COMPILER CONSTRUCTION June 1982, BOSTON, US pages 205 - 215; J. CRAWFORD: 'Engineering a production code generator' * page 213, column 2 - page 214, column 1, line 30 ** - - - -	1	
			TECHNICAL FIELDS SEARCHED (Int. Cl.5)
			G 06 F

The present search report has been drawn up for all claims

Place of search	Date of completion of search	Examiner
The Hague	04 December 91	ADMINISTRATION
CATEGORY OF CITED DOCUMENTS		
<p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : Intermediate document T : theory or principle underlying the invention</p> <p>E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>		